Conditions That Impact the Complexity of QoS Routing

Fernando A. Kuipers and Piet F. A. Van Mieghem, *Member, IEEE*

***Abstract*—Finding a path in a network based on multiple constraints (the MCP problem) is often considered an integral part of quality of service (QoS) routing. QoS routing with constraints on multiple additive measures has been proven to be NP-complete. This proof has dramatically influenced the research community, resulting into the common belief that exact QoS routing is intractable in practice. However, to our knowledge, no one has ever examined which "worst cases" lead to intractability. In fact, the MCP problem is not strong NP-complete, suggesting that in practice an exact QoS routing algorithm may work in polynomial time. The goal of this paper is to argue that in practice QoS routing may be tractable. We will provide properties, an approximate analysis, and simulation results to indicate that NP-completeness hinges on four conditions, namely: 1) the topology; 2) the granularity of link weights; 3) the correlation between link weights; and 4) the constraints. We expect that, in practice, these conditions are manageable and therefore believe that exact QoS routing is tractable in practice.**

***Index Terms*—Complexity, multi-constrained path, QoS routing, phase transition.**

I. Introduction

THERE is an increasing demand for using real-time multimedia applications over the Internet. In order for these applications to work properly, quality of service (QoS) measures like bandwidth, delay, jitter, packet loss, etc., need to be controlled. Currently, the Internet cannot guarantee that the QoS requirements of applications will be satisfied. This has triggered the research community to (en masse) investigate the QoS problem, resulting in proposals for QoS-based frameworks (e.g., IntServ, DiffServ, constraint-based MPLS), QoS routing protocols (e.g., Q-OSPF, PNNI), and many QoS routing algorithms (see [15]).

Routing in general consists of two entities, namely the routing protocol and the routing algorithm. The routing protocol has the task of capturing the state of the network and its available network resources and disseminating this information throughout the network. The routing algorithm uses this information to compute shortest paths. Best-effort routing performs these tasks based on a single measure, usually hop-count. QoS routing, however, must take into account multiple QoS measures and requirements. In this paper, we assume that the network-state information is temporarily static and that it has been distributed throughout the network and is accurately maintained at each node using QoS routing protocols. Once a node acquires the network-state information, it performs the second task in QoS routing, namely computing paths given multiple QoS constraints, also known as the multi-constrained path (MCP) problem. In this paper, we evaluate the complexity of exactly solving the MCP problem. Before giving the formal definition of the MCP problem, let us first describe the notation that is used.

Let $G(N,E)$ denote a network topology, where N is the set of nodes and E is the set of links. With a slight abuse of notation, we also use N and E to denote the number of nodes and the number of links, respectively. The number of QoS measures is denoted by m. Each link is characterized by an m-dimensional link weight vector, consisting of m nonnegative QoS weights $(w_i(u,v), i=1,\ldots,m, (u,v)\in E)$ as components. The QoS measure of a path can either be additive, multiplicative, or min/max. In the case of additive measures (e.g., delay, jitter), the path weight of that measure equals the sum of the QoS weights of the links defining the path. Multiplicative measures can be transformed into additive weights by using the logarithm. The path weight of min(max) QoS measures (e.g., available bandwidth) refers to the minimum(maximum) of the QoS weights along the path. The QoS constraints of an application are expressed in the m-dimensional vector $\vec{L}$. Constraints on min(max) QoS measures can easily be treated by omitting all links (and possibly disconnected nodes), which do not satisfy the requested QoS constraint. In contrast, constraints on additive QoS measures cause more difficulties. Therefore, for our study on complexity, we assume all QoS measures to be additive.

Definition 1: Multi-Constrained Path (MCP) Problem: Consider a network $G(N,E)$. Each link $(u,v)\in E$ is specified by m additive QoS weights $w_i(u,v)\geq 0, i=1,\ldots,m$. Given m constraints $L_i, i=1,\ldots,m$, the problem is to find a path P from a source node s to a destination node d such that

$$w_i(P) \stackrel{\text{def}}{=} \sum_{(u,v)\in P} w_i(u,v) \leq L_i, \quad \text{for } i=1,\ldots,m.$$

There may exist multiple paths in the graph $G(N,E)$ that satisfy all the constraints. Such paths are said to be feasible. According to Definition 1, any of these paths is a solution to the MCP problem. However, it might be desirable to retrieve the optimal path, according to some criterion, within the constraints. This more difficult problem is known as the *Multi-Constrained Optimal Path* (MCOP) problem.

The rest of this paper is organized as follows. Section II presents an overview of related work. Section III analyzes the worst case NP complexity of the MCP problem. Section IV

Manuscript received May 1, 2003; revised February 11, 2004, and September 10, 2004; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor J. Roberts.

The authors are with the Delft University of Technology, Electrical Engineering, Mathematics and Computer Science, 2600 GA Delft, The Netherlands (e-mail: f.a.kuipers@ewi.tudelft.nl; p.vanmieghem@ewi.tudelft.nl).

Digital Object Identifier 10.1109/TNET.2005.852882

evaluates, mathematically and by simulation, the impact of correlation on the complexity of solving the MCP problem. Section V discusses the impact of the constraint values on the complexity and introduces the concept of phase transitions in the MCP problem. Finally, in Section VI, we present our conclusions.

II. Related Work

The MCP problem is an NP-complete problem. Garey and Johnson [8] were the first to list the MCP problem with $m = 2$ as being NP-complete, but they did not provide a proof. Wang and Crowcroft have provided this proof for $m \geq 2$ in [27] and [28], which basically consisted in reducing the MCP problem for $m = 2$ to an instance of the *partition* problem, a well-known NP-complete problem [8]. The effect of this proof has been tremendous, because it suggests that the MCP problem is intractable, in which case heuristics should be used. Many simulations performed in [6], [16], [23], and [25][1] suggest that exact QoS routing may not be intractable in practice. There are certain NP-complete problems, such as *partition*, which are considered by many practitioners to be tractable in practice. The reason for this is that, although no algorithms for solving them in time bounded by a polynomial in the input length (e.g., N, E) are known, there exist algorithms which solve those problems in time bounded by a polynomial in the input length and the magnitude of the largest number (e.g., largest QoS weight) in the given problem instance [9]. Such algorithms are called pseudo-polynomial-time algorithms. NP-complete problems for which no exact pseudo-polynomial-time algorithm exists, are called NP-complete in the strong sense. In the case of the *partition* problem, the NP-completeness strongly depends on the fact that arbitrarily large numbers are allowed. If any upper bound was imposed on these numbers in advance, even a bound which is a polynomial function of the input length, there would exist a polynomial-time algorithm for solving this (restricted) problem [9].

Pisinger [22] has evaluated Knapsack problems, which are NP-complete problems (proved via reduction to the *partition* problem), and found that in practice these problems are tractable. For many more NP-complete problems, typical cases are "easy" to solve. A study of the phenomenon that typical cases are "easy" was performed by Cheeseman *et al.* [4], who introduced the concept of phase transitions in NP-complete problems. According to Cheeseman *et al.*, NP-complete problems which are very under-constrained are soluble and it is usually easy to find one of the many solutions. NP-complete problems which are very over-constrained are insoluble. In the phase transition in between, problems are "critically constrained" and it is typically very hard to determine if they are soluble or insoluble [10]. For a more formal discussion of phase transitions, we refer to [7]. Cheeseman *et al.* have conjectured that all NP-complete problems have at least one order parameter and that the hard to solve problems are around a critical value of this order parameter. Although this conjecture does not hold for all NP-complete problems [13], there seems to be a connection between complexity and phase transitions. The lack of a phase transition seems to have significant computational implications: such problems are either computationally tractable, or well-predicted by a single, trivial algorithm [13]. This alleged connection between complexity and phase transitions motivated us to investigate phase transitions in the MCP problem. Monasson *et al.* [20] report an analytic solution and experimental investigation of the phase transition in K-satisfiability (the first problem shown to be NP-complete). Gent and Walsh [10] show that phase transitions occur in the *partition* problem.

Levin [17] advocated a different study of NP-complete problems by introducing the concept of average-case complexity. He indicated that some NP-complete problems are "easy on average," while other (average-case NP-complete) problems may not be.

There exists also some work in the literature revealing important properties of the MCP problem. We will mention three of those properties, that all strengthen our belief that in practice exact QoS routing is tractable. First, the MCP problem is not strong NP-complete. Second, if all but one measures take bounded integer values, then the MCP problem is solvable in polynomial time [5]. Finally, if some specific dependencies exist between QoS measures, exact QoS routing can be performed in polynomial time [19]. The goal of our work is to provide more evidence that suggests the tractability of exact QoS routing, in practice.

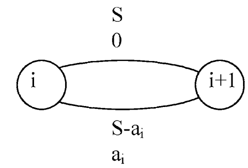

Fig. 1. Assignment of link weights to the links between nodes i and $i+1$, in a chain topology.

III. Worst Case Complexity Analysis

In this section we will analyze the worst case complexity of the MCP problem for $m = 2$. First, we will rewrite the proof that the MCP problem for $m = 2$ is NP-complete [27], [28], and refer to it as the NP-proof.

Theorem 1: The MCP problem is NP-complete.

Proof: First the proof for $m = 2$ is presented. Given a chain topology with $n+1$ nodes and $2n$ links, each with a two-component weight vector $\vec{w}$ as depicted in Fig. 1, and a set of numbers $a_i \in A, 0 \leq a_i \leq S$, for $i = 1, \ldots, n$, where $S = \sum_{i=1}^{n} a_i$. The constraints are chosen as follows: $L_1 = nS - \frac{S}{2}$, and $L_2 = \frac{S}{2}$. To solve the MCP problem, we need to find a path from node 1 to node $n+1$, that obeys the constraints. Since, for all link weight vectors, the sum of the components equals S, we have that $w_1(P) + w_2(P) = nS$. Accordingly, a solution satisfying the constraints is only found if $w_1(P) = nS - \frac{S}{2}$ and $w_2(P) = \frac{S}{2}$. The problem has now become an instance of the well-known NP-complete *partition* problem [8] and can only be solved by finding the set $A' \subseteq A$, for which $\sum_{a_i \in A'} a_i = \frac{S}{2}$. A feasible path exists if the set A' exists, in which case it is retrieved by choosing the lower link if $a_i \in A'$ and the upper link if $a_i \notin A'$.

[1] Our paper [25] can be seen as a more simulative companion to this paper.

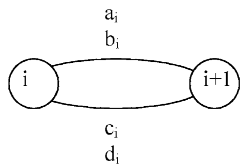

Fig. 2. Chain topology with two QoS weights per link and N nodes in total.

We have proved that the MCP problem with $m = 2$ is NP-complete. The proof that MCP in general is NP-complete inductively follows. We assume that the MCP problem with m measures is NP-complete. If we extend the number of measures with 1 to $m+1$ and choose $L_{m+1} = \sum_{(u,v)\in E} w_{m+1}(u,v)$, then all paths between source and destination obey this constraint. The MCP problem with $m+1$ measures is then only solved if the MCP problem with m measures is solved. This concludes the proof. ■

Corollary 2: The MCP problem is not NP-complete in the strong sense.

Proof: The MCP problem is not strong NP-complete, because there exist pseudo-polynomial algorithms that exactly solve this problem (e.g., see [14] and [18]). ■

The proof that a problem is NP-complete or not is entirely based on a worst case argument. A problem is called polynomially solvable if it can be solved by an algorithm that terminates after a number of steps (instructions) that is bounded by a polynomial in the input length. A problem is called NP-complete if there is even one instance that is not polynomially solvable (unless $P = \mathrm{NP}$). It may occur that in some instances the running time required to solve the MCP problem is polynomial. We call those polynomially solvable instances tractable and we will use the term intractable when instances require a nonpolynomial running time (i.e., they are not polynomially solvable).

We desire to distinguish the instances of the MCP problem that are tractable and those that are intractable. If we look at the graph on which the MCP problem should be solved, we could delineate the class of polynomially solvable graphs, i.e., the class of graphs in which the number of paths between two nodes increases as a polynomial function of N (e.g., tree, circle, and star topologies). This class of graphs is most likely very small and therefore most graphs potentially can lead to intractability. Fortunately, the underlying topology *alone* is not sufficient to lead to intractability: we also need a specific link weight structure. For instance, if all link weights are assigned the value 1, then the MCP problem is polynomially solvable regardless of the underlying topology. We will proceed by defining a link weight structure that leads to intractability in the chain topology. We will use the chain topology as depicted in Fig. 2 and ascertain that all paths from source s to destination d are nondominated.

Definition 2: Dominance: A path P *dominates* a path P' if $w_i(P) \leq w_i(P')$ for all link weight components i except for at least one j for which $w_j(P') < w_j(P)$. A path P is called *nondominated* if there[2] does not exist a path P' for which $w_i(P') \leq w_i(P)$ for all link weight components i except for at least one j for which $w_j(P') < w_j(P)$.

In general, there are two important properties that can reduce the search space when solving the MCP problem without losing exactness, namely nondominance and the constraints themselves. If a sub-path P from source node s to node i exceeds one or more constraints, it can never become a feasible path,[3] because the path weight vector from i to destination node d consists of nonnegative weights. Similarly, if for two paths P_1, P_2 from s to i it holds that P_1 dominates P_2, then all weights of P_1 are smaller than (or equal to) those of P_2 and hence we can omit P_2 from our search space and continue with P_1 [6], because the paths extended from P_2 will always be dominated by the paths extended from P_1. According to [25], the maximum number of nondominated paths that obey the constraints is upper bounded by $\frac{\prod_{i=1}^{m} L_i}{\max_j(L_j)}$, where the constraints L_i are expressed as an integer number of the smallest granularity. This value provides a worst case estimate of the size of our search space. According to Levin [17] some NP-complete problems are "easy on average," while other (average-case NP-complete) problems may not be. The average-case complexity therefore also gives some indication whether an NP-complete problem could be tractable in practice. In [25] we have shown that if the path weights are independently distributed in the solution space, then the MCP problem can be solved in polynomial time on average.

Without loss of generality, we assume that the link weights in Fig. 2 are chosen such that $a_i > c_i$ and $b_i < d_i$, for $i = 1, \ldots, N$ ($c_i > a_i$ and $d_i < b_i$ would also have been possible). It can be verified that if $a_i \geq c_i$ and $b_i \geq d_i$ or $c_i \geq a_i$ and $d_i \geq b_i$ were allowed, this would lead to dominance.

Property 1: If, in the chain topology in Fig. 2, it holds that

$$\begin{cases} a_i - c_i > \sum_{j=0}^{i-1}(a_j - c_j) \\ b_i - d_i < \sum_{j=0}^{i-1}(b_j - d_j) \end{cases} \tag{1}$$

for $i = 1, \ldots, N-1$, where $a_0 = b_0 = c_0 = d_0 = 0$, then all 2^{N-1} paths from node 1 to node N are nondominated.

Proof: We will give a proof by induction.

$i = 1$: There are two paths from node 1 to node 2, namely $P_1(1 \to 2) = \begin{pmatrix} a_1 \\ b_1 \end{pmatrix}$ and $P_2(1 \to 2) = \begin{pmatrix} c_1 \\ d_1 \end{pmatrix}$. According to formula (1): $a_1 > c_1$ and $b_1 < d_1$, which shows that both paths from node 1 to node 2 are nondominated.

The inductive step is to assume the correctness of formula (1) for a certain i. It remains to prove that it also holds for $i+1$. There are 2^{i-1} paths from node 1 to i. From i there are two possible links to $i+1$, resulting in a total of 2^i paths from node 1 to node $i+1$. 2^{i-1} paths will follow the upper link from i to $i+1$, while the remaining 2^{i-1} paths will follow the lower link. Since all paths at i are nondominated (inductive assumption), the paths following the upper link are also nondominated, because the same vector is added to each of the path vectors. The same property applies to the paths that follow the lower link. It remains to show that if (1) holds, then the paths following the

[2]If there are two or more different paths between the same pair of nodes that have an identical weight vector, only one of these paths suffices. In the sequel, we will therefore assume one path out of the set of equal-weight vector paths as being nondominated and regard the others as dominated paths.

[3]This also holds for the lower-bound estimation of the end-to-end path weight vector $\vec{w}(P) + \vec{b}$, where $\vec{b}$ denotes a lower-bounds vector consisting of the m one-dimensional shortest path weights from i to d.

upper link and the paths following the lower link do not dominate each other.

If (1) is satisfied, then all paths following the upper link possess a first path weight larger than the first weights of the paths following the lower link. Similarly, the paths following the lower link have a second weight, which is larger than the second weights of the paths following the upper link. Hence, the paths following different links are nondominated. ■

The *partition* problem is NP-complete, because the values involved in an instance of the *partition* problem may be arbitrarily large (or have an infinite granularity). The same phenomenon is observed in formula (1), where the difference between a_i and c_i (and correspondingly d_i and b_i) grows exponentially:

$$\begin{aligned} a_{i+1} - c_{i+1} &> \sum_{j=0}^{i}(a_j - c_j) = (a_i - c_i) + \sum_{j=0}^{i-1}(a_j - c_j) \\ &> 2\sum_{j=0}^{i-1}(a_j - c_j) > \cdots > 2^{i-1}(a_1 - c_1). \end{aligned}$$

If a_i in the NP-proof are not chosen according to formula (1), but if they take bounded integer values, then the problem becomes polynomially solvable.

A second important phenomenon that we observe from formula (1) is that the link weights display a perfect *negative correlation*. If the link weights would have had a positive correlation, then if $a_i > c_i$ most likely also $b_i > d_i$, leading to dominance.

Lemma 3: Property 1 is a sufficient but also necessary condition for all paths in the chain topology to be nondominated.

Proof: We need to show that if formula (1) does not hold, then at least one path from node 1 to node $i+1$ is dominated. If (1) does not hold, we have

$$\begin{cases} \sum_{j=0}^{i-1} c_j + a_i \le \sum_{j=0}^{i-1} a_j + c_i \\ \sum_{j=0}^{i-1} d_j + b_i \ge \sum_{j=0}^{i-1} b_j + d_i \end{cases} \quad (2)$$

or

$$\begin{cases} \sum_{j=0}^{i-1} c_j + a_i > \sum_{j=0}^{i-1} a_j + c_i \\ \sum_{j=0}^{i-1} d_j + b_i \ge \sum_{j=0}^{i-1} b_j + d_i \end{cases} \quad (3)$$

or

$$\begin{cases} \sum_{j=0}^{i-1} c_j + a_i \le \sum_{j=0}^{i-1} a_j + c_i \\ \sum_{j=0}^{i-1} d_j + b_i < \sum_{j=0}^{i-1} b_j + d_i. \end{cases} \quad (4)$$

We have written these formulas slightly differently from (1) to illustrate that they correspond to two paths, namely the path that followed all the lower links up to node i and took the upper link from node i to node $i+1$ and the path that took all the upper links toward node i and the lower link from node i to node $i+1$. Formula (2), without the equalities, is exactly the same as (1), but a is called c and b is called d. If the equality sign applies, then the path that followed all the lower links up to node i and took the upper link from node i to node $i+1$ is the same as the path that took all the upper links toward node i and the lower link from node i to node $i+1$. According to Definition 1 only one of these two paths is nondominated. When formula (3) applies, the path that followed all the lower links up to node i and took the upper link from node i to node $i+1$ is dominated by (or dominates in the case of formula (4)) the path that took all the upper links toward node i and the lower link from node i to node $i+1$. ■

Property 1 and Lemma 3 seem very restrictive, because they are solely based on the chain topology and we require all paths to be nondominated. If only a subset of all paths (that increases nonpolynomially in N) were nondominated, then the problem would still be intractable. However, if only such a subset of all paths would be nondominated, then Property 1 must hold for a subset of the links/subpaths. Otherwise, all link weights would be bounded and the problem would be polynomially solvable.

Also the chain topology can be put into perspective. Links in the chain topology can be seen as sub-paths.

Lemma 4: If there are more than two links (all with two weights) between two nodes in the chain topology, formula (1) should hold for all possible pairs of links, in order for all paths from node 1 to node N to be nondominated.

In practice we do not expect links/sub-paths to satisfy formula (1). If formula (1) is not satisfied, Lemma 4 suggests that when there are many sub-paths to a node, the probability that all these paths are nondominated decreases and consequently also the search space decreases.

At the beginning of this section we mentioned that there are two important properties to reduce the search space, namely nondominance and the values of the constraints. If the constraints are chosen very large, then it will be easy to find a path that obeys these constraints. On the other hand, if the constraints are very strict, there may not be a path available that can obey these constraints. For the chain topology, besides formula (1), the constraints must lie in the range

$$\begin{cases} \sum_{j=0}^{N-1} c_j \le L_1 \le \sum_{j=0}^{N-1} a_j \\ \sum_{j=0}^{N-1} d_j \ge L_2 \ge \sum_{j=0}^{N-1} b_j \end{cases}$$

to induce intractability. Since $c_i < a_i$, the shortest path for measure 1 from node 1 to node N equals $\sum_{j=0}^{N-1} c_j$. If $L_1 < \sum_{j=0}^{N-1} c_j$, then no feasible path exists. If $L_1 > \sum_{j=0}^{N-1} a_j$, then all possible (loop-free) paths can obey this constraint. The same reasoning applies to L_2 and is further motivated in Section V.

In this section we have used the chain topology to create an intractable instance of the MCP problem. This instance provided us with some hints on the underlying causes of intractability. In Section IV we will further evaluate the impact of correlation on the complexity of QoS routing.

IV. The Impact of Link Correlation on Complexity

Section III hinted at a connection between link correlation and complexity. In this section we will discuss the impact of link correlation on the complexity of QoS routing by giving some properties and presenting simulation results.

A. Theory

Ma and Steenkiste [19] have shown that when specific dependencies (correlation) exist between QoS measures, due to Weighted Fair Queueing scheduling, QoS routing can be performed in polynomial time. However, it is a misconception that if all QoS measures are a function of a common measure, then by just minimizing this common measure, we will have minimized all measures. We will illustrate that this is not always

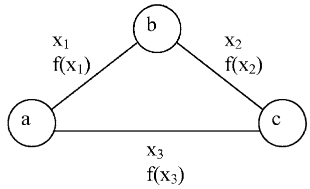

Fig. 3. Example topology.

the case and provide some conditions when this statement holds. We will denote by $f(\cdot)$ a convex function, by $\varphi(\cdot)$ a concave function, by $\psi(\cdot)$ a linear function, and by $g(\cdot)$ a monotone increasing function.

Consider Fig. 3: if $f(x)$ is a convex function, then the shortest path based on x is not necessarily the shortest path for $f(x)$. For example, suppose that $f(x) = e^x$ and $x_1 = 2, x_2 = 2, x_3 = 3$. Then the shortest path from a to c is $a - c$ for x, but $a - b - c$ for $f(x)$.

Likewise, if $\varphi(x)$ is a concave function, the shortest path based on x is not necessarily the shortest path for $\varphi(x)$, e.g., $\varphi(x) = \log(x)$ and $x_1 = 1.2, x_2 = 1.2, x_3 = 2.2$. Then the shortest path from a to c is $a - c$ for x, but $a - b - c$ for $\varphi(x)$.

In case of a linear function $\psi(x) = ax + b$, then the shortest path based on x will also be the shortest path for $\psi(x)$ if $a > 0$ and $b = 0$.

In the rest of this subsection we consider graphs, for which all link weights are a function of a common link weight. Each link i has a weight vector $\vec{w} = \begin{bmatrix} f_1(x_i) \\ \vdots \\ f_m(x_i) \end{bmatrix}$, where x_i is the common link parameter (links may have different x_i and different f_j). In the sequel we will refer to this graph as G_w. We also introduce the graph G_x, which is identical in structure to G_w, but for which the links only have weight x_i.

Let P_x be the shortest path from source s to destination d in G_x, then

$$w(P_x) = \sum_{i \in P_x} x_i \leq w(P) = \sum_{i \in P} x_i$$

where P is any other path ($\neq P_x$) from s to d in G_x. Let $\varphi(x)$ be a concave function, then

$$\varphi\left(\frac{1}{h}\sum_{i=1}^{h} x_i\right) \geq \frac{1}{h}\sum_{i=1}^{h} \varphi(x_i)$$

where h is the hopcount of a path P.

Property 4: If the weight vector of a link, $\vec{w} = \begin{bmatrix} \varphi_1(x_i) \\ \vdots \\ \varphi_m(x_i) \end{bmatrix}$ with $\varphi_j(x_i)$ concave functions, is a function of a single parameter x_i and if P is the shortest path from s to d in G_x with length $X = \sum_{i=1}^{h} x_i$ and hopcount h, then P in G_w satisfies the constraint vector $\vec{L}$ if

$$X \leq h\varphi_j^{-1}\left(\frac{L_j}{h}\right), \quad 1 \leq j \leq m. \tag{5}$$

Proof: The constraints are satisfied if $\sum_{i \in P} \varphi_j(x_i) \leq L_j$. Since φ_j are concave functions

$$\sum_{i=1}^{h} \varphi_j(x_i) \leq h\varphi_j\left(\frac{1}{h}\sum_{i=1}^{h} x_i\right) \leq L_j$$

or

$$\varphi_j\left(\frac{1}{h}\sum_{i=1}^{h} x_i\right) \leq \frac{L_j}{h}.$$

Hence,

$$X = \sum_{i=1}^{h} x_i \leq h\varphi_j^{-1}\left(\frac{L_j}{h}\right).$$

■

Note that although P is the shortest path in G_x, this does not mean that P is also the shortest path in G_w (there may be another path P' for which $\sum_{i \in P'} \varphi(x_i) < \sum_{i \in P} \varphi(x_i)$). Equation (5) is a sufficient but not a necessary condition, because there may be a path that does not obey (5), but still satisfies the constraints.

Property 5: If the weight vector of a link, $\vec{w} = \begin{bmatrix} f_1(x_i) \\ \vdots \\ f_m(x_i) \end{bmatrix}$ with $f_j(x)$ convex functions, is a function of a single parameter x_i and if P is the shortest path from s to d in G_x with length $X = \sum_{i=1}^{h} x_i$ and hopcount h, then P (and therefore all paths) violates the constraints in G_w if

$$X > hf_j^{-1}\left(\frac{L_j}{h}\right) \tag{6}$$

for at least one j.

Proof: By convexity

$$hf_j\left(\frac{1}{h}\sum_{i=1}^{h} x_i\right) = hf_j\left(\frac{X}{h}\right) \leq \sum_{i=1}^{h} f_j(x_i).$$

The jth constraint is violated if $\sum_{i=1}^{h} f_j(x_i) > L_j$, which is the case if $hf_j(\frac{X}{h}) > L_j$, which is equivalent to (6). ■

Property 6: If the weight vector of a link $\vec{w} = \begin{bmatrix} g_1(x_i) \\ \vdots \\ g_m(x_i) \end{bmatrix}$ with $g_j(x_i)$ monotone increasing and P is the shortest minimum-hop path from s to d in G_x and $x_i \leq x_i'$, where x_i' is the ith ordered common link weight of another path P' from s to d in G_x, then P is also the shortest path in G_w.

Proof: The property is a corollary from [11, Th. 107]: Suppose that the sets (a) and (a') are arranged in descending order of magnitude. Then a necessary and sufficient condition that $g(a_1') + \cdots + g(a_n') \leq g(a_1) + \cdots + g(a_n)$ should be true for all continuous and increasing g is that $a_v' \leq a_v (v = 1, 2, \ldots, n)$. ■

B. Simulation Results

In this section we will evaluate the complexity of QoS routing through simulations. We will present simulation results for several classes of graphs, namely the class of random graphs, the class of square two-dimensional (2d) lattices, the class of power-law graphs, and the chain topology. The class of random graphs is of the type $G_p(N)$ [3], where p is the expected link density[4] ($p = 0.2$). For the class of Internet-like power-law graphs, we have chosen the power $\alpha = -2.4$ in the nodal degree distribution $\Pr[d = k] \sim k^{\alpha}$. The chain topologies were of a triangular shape (as depicted in Fig. 3). We have simulated with three different distributions for the $m = 2$ link weights, namely the uniform, exponential, and Gaussian distributions. We only present the simulation results for correlated uniformly distributed link weights $\in [0, 1]$ with correlation coefficient[5] ρ [21], because they led to a higher complexity than the exponential and Gaussian distributions. We have previously also simulated with $m > 2$. The results are scattered over several papers (e.g., [25] and [16]). The results (assuming independence among the m weights) do not show a more than linear increase in complexity as a function of m. We have confined to $m = 2$ for the correlation study, because for $m = 2$ the correlation coefficient can span the entire range $[-1, 1]$, while if m grows, the links cannot all be correlated with $\rho = -1$ and the "mutual" correlation range tends to $[0, 1]$.

All simulations consisted of generating 10^5 different graphs and in each graph a path was computed via the SAMCRA algorithm [26]. SAMCRA incorporates four concepts: 1) a nonlinear measure for the path length $l(P) = \max_{j=1,\ldots m} \left(\frac{w_j(P)}{L_j} \right)$; 2) a k-shortest path approach[6] to examine multiple subpaths per node; 3) the principle of nondominated paths to reduce the search space; and 4) the "look-ahead" concept. The look-ahead concept precomputes one or multiple shortest path trees rooted at the destination and then uses this information to reduce the search space. In TAMCRA [6], the polynomial-time predecessor of SAMCRA, k is fixed (giving its polynomial complexity), but with SAMCRA this k can grow exponentially in the worst case. SAMCRA does not only exactly solve the MCP problem, but also exactly solves the MCOP problem by finding the optimal path within the constraints. Since the MCOP problem is more difficult than the MCP problem, the simulation results presented here should be interpreted as an upper bound. We have simulated a worst case scenario by choosing the constraints so large that all paths can satisfy the constraints. Therefore, SAMCRA must search in the largest search space possible (all nondominated paths between the source and destination), for the optimal path. If SAMCRA was only solving the MCP problem, choosing such large constraints would make the MCP problem "easy," because then any path is a solution to the MCP problem. During all simulations, we kept track of the minimum queue size ($k_{\min}$: the minimum number of paths that needs to be stored at a node) needed to find a feasible path. If TAMCRA [6] had used this particular $k_{\min}$ under the same conditions, it would have found the same optimal path as SAMCRA did. If a smaller queue size had been used, TAMCRA would not have been able to find the optimal path. This minimum queue size $k_{\min}$ can grow as a factorial in the worst case and presents our measure for the complexity of QoS routing.

[4]We have used a fixed link density, because according to [24] the dependency on the link density becomes vanishingly small as N grows.

[5]We have verified that the correlation coefficient ρ' of the generated random variables equals the desired ρ.

[6]A k-shortest path algorithm does not stop when the destination has been reached for the first time, but continues until it has been reached through k different paths succeeding each other in length.

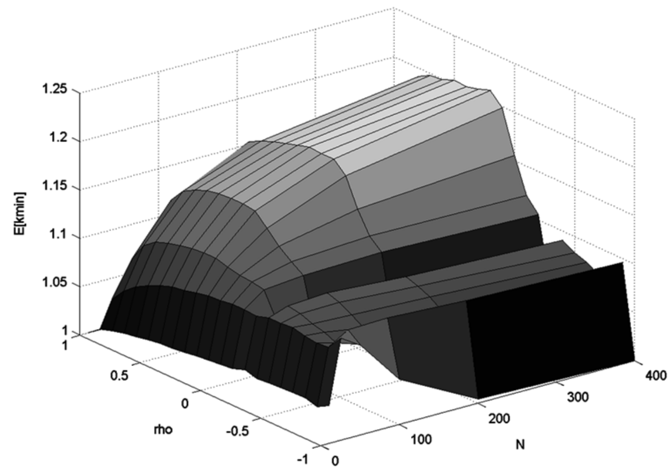

Fig. 4. Expected queue size for the class $G_p(N)$, with $m = 2$ uniformly distributed correlated link weights, as a function of the number of nodes N and the correlation coefficient ρ.

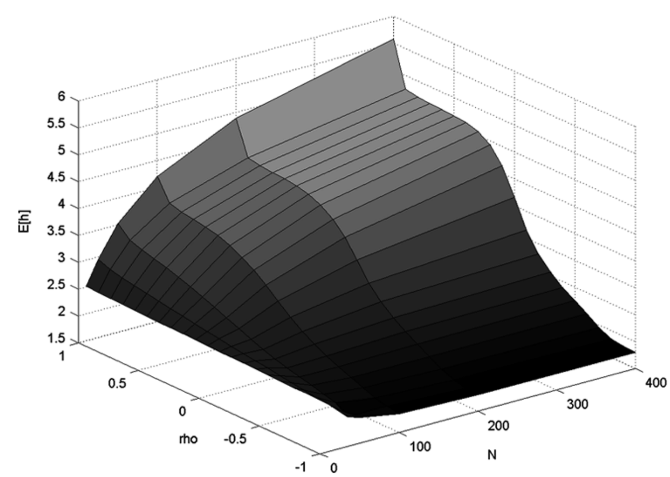

Fig. 5. Expected hopcount for the class $G_p(N)$, with $m = 2$ uniformly distributed correlated link weights, as a function of the number of nodes N and the correlation coefficient ρ.

As illustrated in Fig. 4, the results for the class of random graphs do not display any intractability. We can see that a positive correlation leads to a slightly higher $E[k_{\min}]$ than with a negative correlation. This peculiar phenomenon has only been observed in the class of random graphs, with correlated uniformly distributed link weights. An explanation can be found if we look at Fig. 5, which shows that a positive correlation between the link weights may induce a higher expected hopcount. When the link weights become more positively correlated, the weights become similar, and the problem approaches the $m = 1$ case. Since the expected hopcount of the m-dimensional shortest paths approaches the minimum hopcount if m grows to infinity [23], the $m = 1$ case is expected to have the largest hopcount. A negative correlation between the link weights also leads to shorter hopcount paths. A low hopcount is possible because there are sufficiently many paths in $G_p(N)$, which

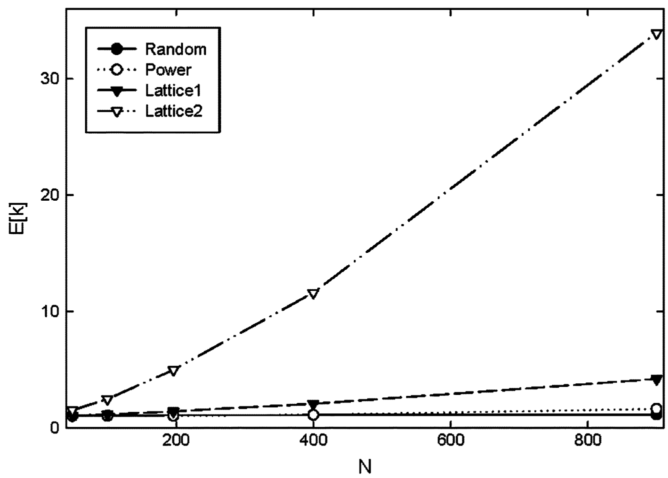

Fig. 6. Expected queue size for different topology classes as a function of the number of nodes N, with $m = 2$ independent ($\rho = 0$) uniformly distributed link weights.

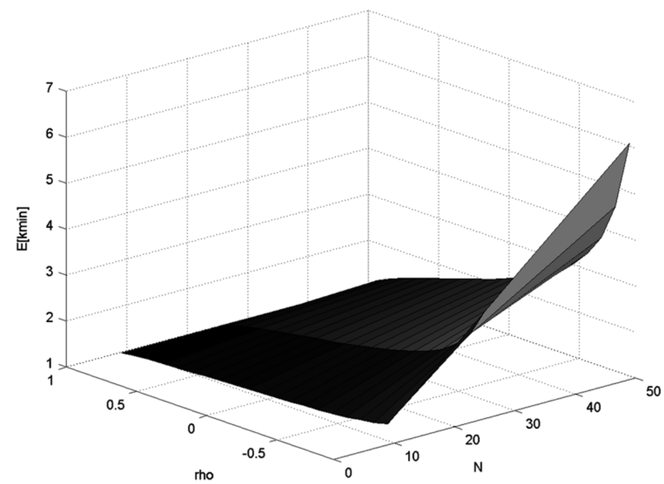

Fig. 7. Expected queue size in the class of two-dimensional lattices as a function of the number of nodes N and correlation coefficient ρ. The $m = 2$ link weights were uniformly distributed and the source and destination nodes were chosen in opposite corners.

can be viewed as a thinning of a complete graph provided $p > \frac{\ln N}{N}$. For negative correlated link weights, a small link weight component is likely accompanied with a large one. For perfect negatively correlated link weight components ($\rho = -1$), SAMCRA's shortest-length path (15) compensates outliers in the link weight components with the result that (one or two) links with weight components close to $\frac{1}{2}$ are selected which leads to the observed minimum-hop paths.

In general, the more hops we must traverse to find the shortest path, the more (sub)-paths we must evaluate and the more complex the computation becomes. We believe that one of the measures for the "computational complexity" of a class of topologies is the expected (minimum) hopcount of an arbitrary path in that topology. The expected hopcount (for $m = 1$) scales as $O(\log N)$ in a random graph, while as $O(\sqrt{N})$ in a 2d lattice and $O(N)$ in the chain topology. Besides the expected hopcount in a graph, also the number of paths between a source and destination can provide a measure for the "computational complexity" of a class of topologies. The class of random graphs with $p = 0.2$ and N increasing, has an increasing number of paths and an increasing average nodal degree, giving the graph a small diameter (i.e., the source and destination are directly linked or a few hops apart). This can be interpreted from Fig. 5. Fig. 6 gives the expected queue size for three different classes of graphs, namely the random graphs ($p = 0.2$), the 2d lattices, and the Internet-like power-law graphs (with power $\alpha = -2.4$). For all three classes of graphs, the source and destination nodes were chosen randomly. Only for the class of 2d lattices "Lattice2," we have chosen the source and destination nodes in opposite corners, to attain the largest minimum hopcount. In the class of random graphs $G_p(N)$, although the number of paths is large, the expected hopcount is small, leading to a small complexity. For the extreme regular class Lattice2 of 2d lattices, the number of paths and the expected hopcount are large, which leads to a large complexity. The class of power-law graphs may be considered, in terms of randomness, to lie between the random graphs and the 2d lattices. The power-law graphs with $\alpha = -2.4$ have a moderate expected hopcount and a small number of paths, and lie, in terms of complexity, closer to the class of random graphs than to the class of 2d lattices. We have also simulated with different link weight distributions, namely Gaussian and exponentially distributed correlated link weights. If we use exponentially distributed correlated link weights, the first weight has a higher probability of being small, than with a uniform distribution. With a uniform distribution, each value for the first weight is equiprobable. Therefore, with exponentially (and also Gaussian) distributed correlated link weights, there is a higher probability that the link weight vectors are similar. For uniformly distributed link weights there is a larger variability, leading to a somewhat worse performance than in the exponential (or Gaussian) case. However, in all cases the expected queue size in the class of random graphs was close to one, leading to a complexity similar to that of Dijkstra's algorithm. These simulation results therefore suggest that, irrespective of the link weight structure, QoS routing in the class of random graphs (and according to [24] also Waxman graphs) is possible in polynomial time. In contrast, the regularity and large expected hopcount in the class of 2d lattices, may provide ground for intractability. Indeed, we can observe a tendency toward intractability in Fig. 7 and true nonpolynomial behavior in Fig. 8.

Because the chain topology was used in the proof that the MCP problem is NP-complete, we have also evaluated the performance of SAMCRA in chain topologies. The results are plotted in Figs. 8 and 9.

Our simulation results[7] indicate that in the class of 2d lattices and chain topologies, the MCP problem seems tractable for nearly the entire range of correlation coefficient ρ, except for extreme negative values. Recall that the NP-proof is based on an extreme negative link correlation. We doubt that in practice link weights will display such a negative correlation, suggesting that exact QoS routing in practice, irrespective of the underlying topology, is possible in polynomial time.

[7]Recall that the simulation results reflect the complexity of the much more difficult MCOP problem.

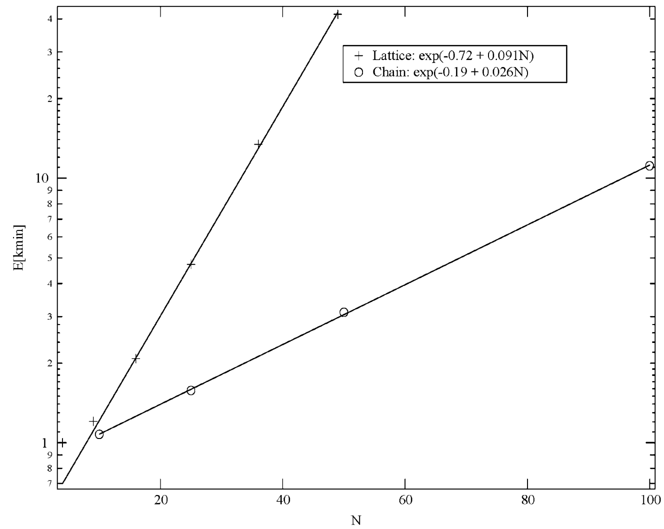

Fig. 8. Expected queue size (on a logarithmic scale) in the class of two-dimensional lattices and chains, as a function of the number of nodes N, with correlation coefficient $\rho = -1$. The $m = 2$ link weights were uniformly distributed and the source and destination nodes were chosen in a way that the minimum hopcount was largest. We have fitted with exponentials, which perfectly match the results in the simulated range. Simulating with larger N may consume months of CPU time and therefore can only be done by reducing the number of iterations or via parallel procesing.

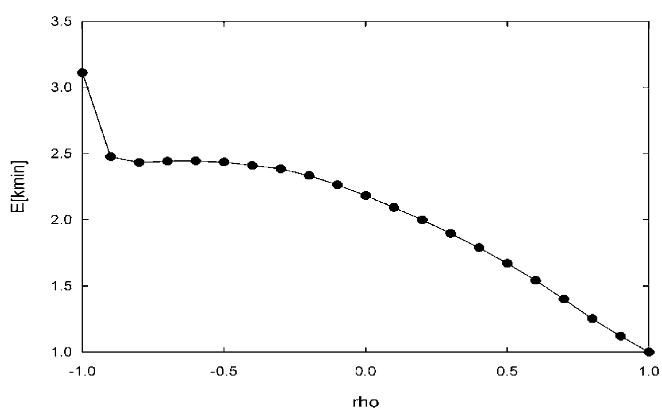

Fig. 9. Expected queue size in the chain topology, with $m = 2$ correlated uniformly distributed link weights for $N = 50$, as a function of the correlation coefficient ρ.

V. The Impact of Constraints on Complexity

In this section we analyze the influence of the constraints on the complexity of the MCP problem. For this purpose, we will initiate an evaluation of a phase transition [4], [12] in the MCP problem.

A. Theory

Property 7: Let $P_{s-d;i}$ denote the one-dimensional shortest path from source s to destination d, for which $w_i(P_{s-d;i}) \leq w_i(P^*)\ \forall P^*$. Then, the MCP and MCOP problems are not NP-complete when

$$L_i < w_i(P_{s-d;i}) \tag{7}$$

for at least one constraint.

Proof: $P_{s-d;i}$ is the path with the shortest ith weight $w_i(P_{s-d;i})$. Therefore, $w_i(P_{s-d;i})$ is a lower bound on the ith weight $w_i(P_{s-d})$ that any path P_{s-d} between s and d can attain. Therefore, if for any constraint i it holds that $L_i < w_i(P_{s-d;i})$, then no path P_{s-d} can obey L_i. Since $P_{s-d;i}$ can be found in polynomial time (e.g., via the Dijkstra algorithm), the MCP problem is solvable (i.e., it is verified that no solution exists) in polynomial time if any constraint obeys (7). ■

Property 8: Let $P_{s-d;i}$ denote the one-dimensional shortest path from source s to destination d for which $w_i(P_{s-d;i}) \leq w_i(P^*)\ \forall P^*$. Then, the MCP problem is not NP-complete when

$$L_i \geq \max_{j=1,\ldots,m}(w_i(P_{s-d;j})) \tag{8}$$

for at least $m-1$ constraints.

Proof: If $L_i \geq \max_{j=1,\ldots,m}(w_i(P_{s-d;j}))$ for all m constraints, then all m one-dimensional shortest paths $P_{s-d;i}$, (for $i = 1,\ldots,m$) obey the constraints. Hence, any path $P_{s-d;i}$ can be chosen as a feasible path.

If $L_i \geq \max_{j=1,\ldots,m}(w_i(P_{s-d;j}))$ for $m-1$ constraints (say $i = 1,\ldots,m-1$) and $L_i < \max_{j=1,\ldots,m}(w_i(P_{s-d;j}))$ for one constraint ($i = m$), then if $L_m \geq w_m(P_{s-d;m})$ path $P_{s-d;m}$ obeys all m constraints. If $L_m < w_m(P_{s-d;m})$, then by Property 7 we know that no feasible path exists. Since the paths $P_{s-d;i}$ can be found in polynomial time (e.g., via the Dijkstra algorithm), the MCP problem is solvable in polynomial time if at least $m-1$ constraints obey (8). ■

For $m = 2$, Properties 7 and 8 constitute a closed NP-complete range

$$w_i(P_{s-d;i}) < L_i < \max_{j=1,\ldots,m}(w_i(P_{s-d;j})). \tag{9}$$

The MCP problem with $m = 2$ is only NP-complete if both constraints lie in the NP-complete range (9). When the link weights are positively correlated, the NP-complete range (9) will be smaller than when the link weights are negatively correlated. This is illustrated in Fig. 10, for $m = 2$. At the cost of increased (polynomial-time) complexity, we can further reduce the NP-complete range by using Property 9.[8]

Property 9: Let P_{s-d} denote the path from source s to destination d for which $\sum_{i=1}^m \alpha_i w_i(P_{s-d}) \leq \sum_{i=1}^m \alpha_i w_i(P^*_{s-d})\quad \forall P^*_{s-d}$. Then, if

$$\sum_{i=1}^m \alpha_i L_i < \sum_{i=1}^m \alpha_i w_i(P_{s-d})$$

where $\alpha_i \geq 0$ with an inequality for at least one i, then there is no feasible path present that can solve the MCP or MCOP problem.

Proof: A proof by contradiction. Assume that P_{s-d} denotes the path from source s to destination d for which $\sum_{i=1}^m \alpha_i w_i(P_{s-d}) \leq \sum_{i=1}^m \alpha_i w_i(P^*_{s-d})\ \forall P^*_{s-d}$ and that $\sum_{i=1}^m \alpha_i L_i < \sum_{i=1}^m \alpha_i w_i(P_{s-d})$. If a path P^*_{s-d} existed that obeys the constraints, then $\sum_{i=1}^m \alpha_i w_i(P^*_{s-d}) \leq \sum_{i=1}^m \alpha_i L_i$, for $i = 1,\ldots,m$ and consequently $\sum_{i=1}^m \alpha_i w_i(P^*_{s-d}) \leq \sum_{i=1}^m \alpha_i L_i < \sum_{i=1}^m \alpha_i w_i(P_{s-d})$, which contradicts our assumption that $\sum_{i=1}^m \alpha_i w_i(P_{s-d}) \leq \sum_{i=1}^m \alpha_i w_i(P^*_{s-d})\ \forall P^*_{s-d}$. Since the path P_{s-d} can be found in polynomial time (e.g.,

[8]We have not programmed Property 9 in our simulations.

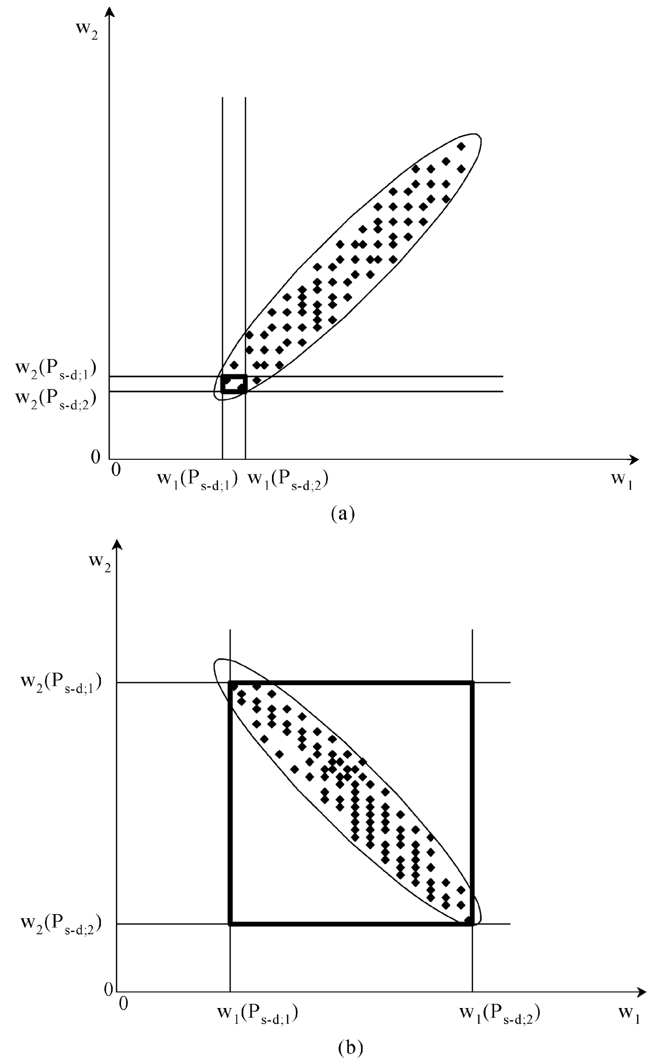

Fig. 10. Constraints range (bold square) for (a) positive correlation and (b) negative correlation. The dots in the figure denote paths in the two-dimensional space ($m = 2$).

via the Jaffe algorithm [14]), the MCP problem is solvable in polynomial time if $\sum_{i=1}^{m} \alpha_i L_i < \sum_{i=1}^{m} \alpha_i w_i(P_{s-d})$. ■

The work presented in Section II suggested that there is a connection between worst case complexity and phase transitions. Using the terminology of Gent and Walsh [10], if problems are very under-constrained, then it is usually easy to find one of the many solutions. When problems are very over-constrained, it is usually easy to determine that they are insoluble. In the phase transition in between, problems are "critically constrained" and it is typically very hard to determine if they are soluble or insoluble. Applied to the MCP problem, we can distinct a phase transition based on the values of the constraints. If one of the constraints obeys (7), the probability of finding a path obeying the constraints is zero. Moreover, it can be verified in polynomial time, that there exists no path in the graph that obeys the constraints (Property 7). On the other hand, if the values of the constraints are very large (under-constrained), such that all constraints follow (8), then a path satisfying these large constraints can be found in polynomial time. A phase transition is therefore expected to occur if the constraints do not obey (7) and (8). For small values of $L_i = w_i(P_{s-d,i}) + \epsilon$ (with $\epsilon > 0$) the MCP problem may still be insoluble, however the effort (complexity) needed to verify that indeed no feasible path is present in the graph has increased. In contrast to the case where the constraints $L_i < w_i(P_{s-d,i})$, only computing the m Dijkstra shortest paths is not sufficient to determine that the problem is insoluble. The SAMCRA [26] algorithm (or another exact MCP routing algorithm) must be invoked and will eventually observe that no path can obey the constraints. The larger the constraints become, the longer it will take to determine that no feasible path exists. Hence, increasing the constraints until a feasible path emerges augments the complexity of its solution. On the other hand, when decreasing the constraints starting from the upper boundary (8), first many paths will obey the constraints $L_i = \max_j(w_i(P_{s-d,j})) - \epsilon$ leading to a high probability that a feasible path will be found fast. If the values of the constraints decrease, the probability of finding a feasible path fast will also decrease. It is therefore expected that a phase transition occurs if there are only a few (if any) feasible paths present. In this case MCP $\approx$ MCOP. The steepness of the phase transition depends on the range between (7) and (8), which is heavily influenced by the correlation coefficient ρ as illustrated in Fig. 10 (and by the computations in the Appendix). As discussed in Section IV, the correlation coefficient also impacts the level of complexity, which decreases if ρ increases.

B. Simulation Results

To be able to observe a phase transition, we must choose an intractable configuration. The simulation results in the previous section suggest that the graphs should contain many paths, have a large expected hopcount, and the link weights should have a negative correlation. All these properties are present in the class of 2d lattices, which in terms of structure and complexity can be seen as a counterpart of the class of random graphs. In the remainder of this paper we confine attention to this class of lattices and try to distinguish a phase transition via simulations and an approximate analysis. For our simulations, we have chosen to use a single 2d lattice with $N = 49$ nodes and correlated uniformly distributed link weights in the range [0, 1].

A worst case scenario is obtained if the source node is positioned in the upper left corner and the destination node in the lower right corner, causing the largest minimum hopcount. For each constraint L_1 and L_2, 100 different values were chosen in the NP-complete range (9) as discussed above, leading to a total of 10^4 iterations, all in the same lattice. Fig. 11 displays the maximum queue size k used by SAMCRA,[9] for $N = 49$ and $\rho = -1$.

Different constraints can lead to different m-dimensional shortest paths. For instance, if L_1 is small (e.g., 5.0 in Fig. 11) and L_2 is large (e.g., 7.0 in Fig. 11), then a path P obeying these constraints must also have a small weight $w_1(P) \leq L_1$ and the second weight may be large as long as $w_2(P) \leq L_2$. Since L_1 is slightly larger than the weight $w_1(P_{s-d;1})$ of the shortest Dijkstra path for measure 1, the path P may closely

[9] k is different from the previously used $k_{\min}$, since k denotes the maximum queue size in SAMCRA whereas $k_{\min}$ is the queue size that TAMCRA would have needed to attain the same solution as SAMCRA. We have used this larger value here, because $k_{\min} = 0$ if there is no path present.

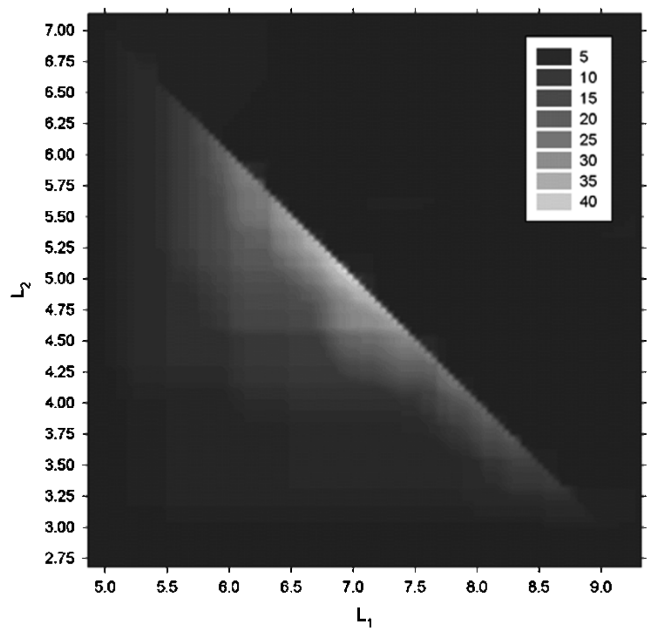

Fig. 11. Contour plot of the queue size in a two-dimensional lattice, with correlated uniformly distributed link weights, $N = 49$, $\rho = -1$, and 10^4 different constraint vectors.

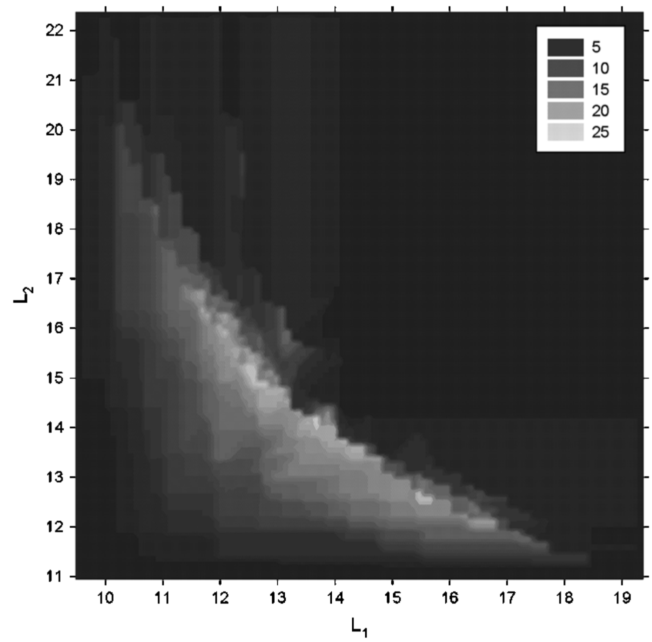

Fig. 12. Contour plot of the queue size in a two-dimensional lattice, with uniformly distributed link weights, $N = 400, \rho = 0$, and 10^4 different constraint vectors.

approximate $P_{s-d;1}$, which may be easy to find as indicated by small k values in Fig. 11. Similarly, if L_1 is large (e.g., 9.0 in Fig. 11) and L_2 is small (e.g., 3.0 in Fig. 11), then a path P obeying these constraints may closely approximate the Dijkstra shortest path for measure 2 $P_{s-d;2}$, which may also be easy to find (as verified in Fig. 11). We observe that the complexity is largest when $L_1 = 6.94$ and $L_2 = 5.06$. These values are situated near the center of the rectangle (Fig. 10) spanned by the NP-complete range (9) at $L_1^* = 7.09$ and $L_2^* = 4.91$. These observations seem to suggest that the complexity is largest when the constraints closely approximate the weights of the m-dimensional shortest path P, which equal $\sqrt{N} - 1$ on average [see the Appendix, (19)]. For 2d lattices of $N = 49$ nodes, we therefore expect the highest complexity for $L_1 = L_2 = 6$. The deviation in our case is caused by only examining one single lattice, instead of the many required for statistical results.

The sharp edge/line in Fig. 11, constituted by the different shortest paths, can be attributed to the extreme negative correlation ($\rho = -1$) as explained in Fig. 10(b) and the Appendix. Since the link weights are chosen in the range [0, 1], we have that for $\rho = -1, w_1(u,v) = 1 - w_2(u,v)\ \forall (u,v) \in E$. Hence, the path weights of any path P obey $w_1(P) = h - w_2(P)$, where $w_i(P) = \sum_{(u,v)\in P} w_i(u,v)$ and h equals the hopcount of path P. If we again look at Fig. 11, we may observe that the straight line, once continued, intersects both axes L_1 and L_2 at 12, which is precisely the minimum hopcount of the 2d lattice with 49 nodes. Moreover, since $w_1(P) = h - w_2(P)$, we know (see Property 8) that when $L_1 + L_2 < h$, then no feasible path exists. This means that for the class of 2d lattices with correlated ($\rho = -1$) uniformly distributed link weights, the constraints must obey $L_1 + L_2 \geq h$, for a feasible path to be possible. This condition for the constraints can be checked in polynomial time and it is therefore possible to obtain a much steeper phase transition than observed in Fig. 11. Finally, we have also simulated with independent uniformly distributed link weights ($\rho = 0$) in the range [0, 1]. As discussed in Section IV, the complexity of solving the MCP and MCOP problems under independent link weights is smaller than with negatively correlated link weights. To observe a phase transition, we had to simulate with a lattice larger than $N = 49$. Fig. 12 gives the contour plot for $N = 400$ and $\rho = 0$. The complexity is largest for $L_1 = 12.58$ and $L_2 = 15.11$.

It would be desirable to obtain an estimation of the size of the constraints that make the MCP problem critically constrained. Such an estimation would allow us to predict the location of the phase transition and hence give us an indication of the "critically constrained" region. In the next subsection we will attempt to provide an approximate analysis of the weights of the m-dimensional shortest path, because as we have seen above, choosing the constraints close to these weights may lead to a nonpolynomial running time.

C. Estimation of the Length of the Shortest Path in a Lattice

This last subsection discusses the approximate computation of the length of the m-dimensional shortest path between two corner points in a rectangular 2d lattice with z_1 links vertically and z_2 links horizontally. The link weights are independent uniformly distributed in the range (0, 1]. The approximate analysis of the formulas presented in this subsection and some of the notation that is used can be found in the Appendix. The asymptotic average weight of a $h = z_1 + z_2$ hop path in one dimension for a 2d lattice is given by (13) as $E[W_1] \simeq \frac{h}{2e} \approx \frac{\sqrt{N}}{e}$. This estimate agrees reasonably well with simulations in the range $N \in [100, 1600]$, which accurately follow $E[W_{\mathrm{sim}}] \approx 0.6N^{0.48}$.

The extension to m dimensions with independent link weight components ($\rho = 0$) for the average length $W_m = L_{\text{eq}} l_h$ is the approximation (17)

$$E[W_m] \simeq \frac{h}{e2^{1/m}}.$$

The scaling $2^{-\frac{1}{m}}$ as a function of m has been observed in simulations, even for $N = 49$. This approximate analysis (16) shows that there is no shortest path obeying the constraints if the length, as defined in (15), $l_h(P) > 1$. This event has probability

$$\Pr[l_h > 1] \approx \exp\left(-\frac{h!}{z_1! z_2!}\left(\frac{L_{\text{eq}}^h}{h!}\right)^m\right).$$

Clearly, if the lattice (i.e., z_1, z_2, and $h = z_1 + z_2$) is fixed and the constraints decrease (increase), all (no) paths violate the constraints. The fact that there exists a path within the constraints depends on the product of the constraints or equivalent constraint L_{eq}. If $\frac{L_{\text{eq}}^h}{h!} > 1$ or $L_{\text{eq}} > (h!)^{\frac{1}{h}} \approx \frac{h}{e}$ (for large h), nearly all paths obey the constraints. If $\frac{L_{\text{eq}}^h}{h!} < 1$ or $L_{\text{eq}} < (h!)^{\frac{1}{h}} \approx \frac{h}{e}$, for a large number m of constraints, no path obeys the constraints. Hence, for large m and large h, there seems to be a critical value of the equivalent constraint $L_{\text{eq}} > (h!)^{\frac{1}{h}} \approx \frac{h}{e}$ for which $E[l_h] = (\frac{z_1! z_2!}{h!})^{\frac{1}{mh}} < 1$ and specifically for the square lattice $E[l_h] \approx 2^{-\frac{1}{m}}$. Below that value the shortest path behavior is clearly different than above that value, which points to a phase transition.

The result (18) in two dimensions ($m = 2$), with perfectly negative correlation ($\rho = -1$), even points to a more confining situation, as was readily observed by comparing Figs. 11 and 12. Since $E[L_{\text{eq}} l_h] \approx \frac{h}{2}$ [see (19)] and any random variable $L_{\text{eq}} l_h \geq \frac{h}{2}$, the average weight of the shortest path lies very close to the boundary $\frac{h}{2}$.

In summary, we have estimated the average length or weights of the shortest path for large values of h or, equivalently, the number of nodes N in the 2d lattice. As common for extremal distributions, the variance is small, which implies a fast transition from 0 to 1 of $\Pr[L_{\text{eq}} l_h \leq y]$ around the average. The knowledge of the shortest path is important to set the constraints: if the constraints are close to $E[L_{\text{eq}} l_h]$, the problem is critically constrained and more computations are needed to determine whether there exists a path obeying the constraints or not. For constraints larger or smaller than $E[L_{\text{eq}} l_h]$, the problem is either under- or over-constrained and the verdict that there exists a path within the constraints is usually simple to draw with high probability. In the analysis presented in the Appendix, we have assumed that a possible overlap of h-hop paths is sufficiently weak to allow the application of the limit laws for independent random variables. Only relatively few paths will share a large number of links. We have used a heuristic argument to validate this assumption and have observed a good agreement with our simulation results. The second assumption is that the shortest path in the 2d lattice has h hops or that $\Pr[\text{hops} > h]$ is negligibly small. This approximation is reasonable since simulations show that $\Pr[\text{hops} = h + 2k]$ is rapidly decaying in k with decay rate dependent on the size of the graph. The larger the graph, the slower the decay rate. However, for increasing m, simulations show that the shortest path tends to have h hops. Also for very negative correlation coefficients, the probability that shortest paths have h hops increases. Finally, although computed for uniformly distributed link weights, the same results hold for any distribution whose h-fold convolved distribution also behaves as $\frac{x^h}{h!}$ for small x. Any distribution in the same sphere of minimal attraction (such as exponentially distributed link weights with mean 1) yields the same results.

VI. Conclusion

In this paper, we have evaluated the complexity of QoS routing. Finding a path based on multiple QoS constraints is proven to be an NP-complete problem. However, this Multi-Constrained Path (MCP) selection problem is not NP-complete in the strong sense, meaning that a pseudo-polynomial algorithm can exactly solve the problem. The NP-completeness of the MCP problem hinges on four factors, namely: 1) the underlying topology; 2) link weights that can grow arbitrarily large or have an infinite granularity; 3) a very negative correlation among the link weights; and 4) the values of the constraints. If the values of the constraints are very large then it is easy to find a path within the constraints. On the contrary, if the values of the constraints are very small, then it is easy to verify that there is no path within the constraints. This indicates that there will be a phase transition if the constraints are around the weights of the m-dimensional shortest path in the network. In this case, it is expected to be difficult to establish whether a feasible path exists. If the four above-mentioned conditions are all necessary to induce intractability, they will allow network and service providers to properly dimension their network and to avoid intractable scenarios. Moreover, if the theory of phase transition holds for the MCP problem, then we know that QoS requirements close to the m-dimensional shortest path will, if admitted, provide the highest possible level of QoS, but also the highest computational cost. Such information is invaluable for pricing and billing mechanisms and admission control algorithms. Finally, a proper understanding and use of the four conditions, will allow for efficient QoS routing at controlled computational costs.

Appendix

In this Appendix, we will present an approximate analysis of the length of the m-dimensional shortest path in a 2d lattice.

A. Analysis for a Single Link Weight ($m = 1$)

Consider a rectangular 2d lattice with size z_1 and z_2 and with independent uniformly distributed link weights on $(0, 1]$. The shortest hop path between two diagonal corner points consists of $h = z_1 + z_2$ hops. The weight W_h of such a h-hop path is the sum of h independent uniform random variables u_j and $W_h = \sum_{j=1}^{h} u_j$ has distribution

$$F(x) = \Pr[W_h \leq x] = \frac{1}{h!} \sum_{j=0}^{h} \binom{h}{j} (-1)^j (x - j)^h 1_{j \leq x}. \tag{10}$$

In particular, $\Pr[W_h \le h] = 1$ and for small $x < 1$ holds that $F(x) = \frac{x^h}{h!}$. We assume that the number[10] $l = \binom{z_1 + z_2}{z_1} = \frac{h!}{z_1! z_2!}$ of those h-hop paths is large. Although these paths can possibly overlap, we ignore this dependence for the moment and assume that the minimum weight among all h-hop paths is well approximated by the limit law (of extremal types [2]) for the minimum of a set of independent random variables X_k with identical distribution F. In particular, if $\lim_{l\to\infty} l(F(x_l)) = \zeta$

$$\lim_{l\to\infty} \Pr[\min_{1\le k\le l} X_k > x_l] = e^{-\zeta}. \tag{11}$$

The limit sequence must obey $l(F(x_l)) \to \zeta$ for sufficiently large l, which implies that $F(x_l)$ must be small or, equivalently, x_l must be small. Hence, $l\frac{x_l^h}{h!} = \zeta$ or $x_l = (\frac{h!\zeta}{l})^{\frac{1}{h}}$. The limit law (11) for the minimum weight $W = \min_{1\le k\le l} W_{h,k}$ of the shortest hop path between two corner points in a rectangular 2d lattice is

$$\lim_{l\to\infty} \Pr\left[\min_{1\le k\le l} W_{h,k} > \left(\frac{h!x}{l}\right)^{\frac{1}{h}}\right] = e^{-x}.$$

In other words, the random variable $\frac{lW^h}{h!}$ tends to an exponential random variable with mean 1 for large $l = \frac{h!}{z_1!z_2!}$ or

$$\Pr[W \le y] \approx 1 - \exp\left(-\frac{y^h}{z_1!z_2!}\right).$$

The mean shortest weight of a h-hop path equals

$$\begin{aligned} E[W] &= \int_0^\infty (1 - F_W(x))dx \approx \int_0^\infty \exp\left(-\frac{x^h}{z_1!z_2!}\right)dx \\ &= \Gamma\left(1 + \frac{1}{h}\right)(z_1!z_2!)^{\frac{1}{h}}. \end{aligned} \tag{12}$$

For a square 2d lattice where $z_1 = z_2 = \frac{h}{2}$, we have

$$E[W] = \Gamma\left(1 + \frac{1}{h}\right)\left(\left(\frac{h}{2}\right)!\right)^{\frac{2}{h}}.$$

Using Stirling's formula ([1, 6.1.38]) for the factorial $h! = \sqrt{2\pi}h^{h+\frac{1}{2}}e^{-h+\frac{\theta}{12h}}$ where $0 < \theta < 1$, we finally arrive for large h at

$$E[W] \simeq \left(\frac{h}{2e}\right)\left(\sqrt{\pi h}e^{\frac{\theta}{6h}}\right)^{\frac{2}{h}} \approx \frac{h}{2e}. \tag{13}$$

We now provide a heuristic argument why, for large h, the neglect of the dependence between h-hop paths is justified. Denote by Γ_h the set of all h-hop paths in the 2d lattice between corner points, with the number of those paths $|\Gamma_h| = \binom{h}{z_1}$. A particular path of the set Γ_h is denoted by γ_h. We denote the weight of γ by $w(\gamma)$. Let w_N be the (random) weight of the shortest path between corner points in the 2d lattice with independent uniformly distributed link weights. The event $\{h_N = h, w_N \le z\}$ implies that there is a h-hop path γ_h with weight $w(\gamma_h) \le z$ and, therefore,

$$\begin{aligned} \Pr[h_N = h, w_N \le z] &\le \Pr[\cup_{\gamma\in\Gamma_h}\{w(\gamma) \le z\}] \\ &\le \sum_{\text{all }\gamma} \Pr[\gamma \in \Gamma_h, w(\gamma) \le z] \end{aligned} \tag{14}$$

where the second inequality follows from Boole's inequality ($\Pr[\cup A_j] \le \sum \Pr[A_j]$). Using the independence of the link and the link weights

$$\begin{aligned} \Pr[h_N = h, w_N \le z] &\le \sum_{\text{all }\gamma} \Pr[\gamma \in \Gamma_h]\Pr[w(\gamma) \le z] \\ &= E[|\Gamma_h|]\Pr[w(\gamma_h) \le z] \end{aligned}$$

or since $\Pr[w(\gamma_h) \le z] = \Pr[W_h \le z]$ given by (10)

$$\Pr[h_N = h, w_N \le z] \le \binom{h}{z_1}F(z).$$

From this rigorous inequality we infer the heuristic argument $\Pr[h_N = h, w_N \le z] \simeq (\genfrac{}{}{0pt}{}{h}{z_1})F(z)$. For a typical value of z, the probabilities should sum to 1, yielding

$$1 = \sum_{j=0}^{\infty} \Pr[h_N = h + 2j, w_N \le z] \simeq F(z)\binom{h}{z_1}$$

where the assumption is that $\sum_{j=1}^{\infty} \Pr[h_N = h + 2j, w_N \le z] \ll \Pr[h_N = h, w_N \le z]$. Hence, a *typical* value for the weight of the shortest path is the solution of $F(z) = \frac{1}{\binom{h}{z_1}}$.

For small z, we have $F(z) = \frac{z^h}{h!}$ such that

$$z \sim (z_1!z_2!)^{\frac{1}{h}}$$

which agrees with $E[W]$ in (12).

B. Analysis for Multiple Link Weights $(m > 1)$

Let us now consider a 2d lattice where each link is specified by a link weight vector $\vec{w} = (w_1, w_2, \ldots, w_m)$. We further confine to the case where all link weight components are independent and uniformly distributed. Using the nonlinear length of SAMCRA [26], the length of a h-hop path is computed as

$$l_h(P) = \max_{1\le j\le m}\left[\frac{W_{h,j}}{L_j}\right] \tag{15}$$

where each weight per component j is $W_{h,j} = \sum_{n=1}^{h} u_{n,j}$ with distribution F given in (10). Since all link weight components are independent

$$\Pr[l_h(P) \le x] = \prod_{j=1}^{m} F(L_j x).$$

For small x

$$\prod_{j=1}^{m} F(L_j x) \approx \prod_{j=1}^{m} \frac{(L_j x)^h}{h!} = \left(\frac{x^h}{h!}\right)^m \prod_{j=1}^{m} L_j^h.$$

We define an equivalent constraint $L_{\text{eq}} = (\prod_{j=1}^{m} L_j)^{\frac{1}{m}}$. Neglecting the dependence of h-hop paths due to possible overlap

[10] Any path in a rectangular lattice can be represented by a sequence of r(ight), l(eft), u(p), and d(own). A shortest hop path between diagonal corner points consists of z_1 r's (or l's) and z_2 d's (or u's). The total number of these paths equals $\binom{z_1 + z_2}{z_1}$.

as above and applying the limit law for the minimum length with $\lim_{l\to\infty} l(\prod_{j=1}^{m} F(L_j x_l)) = \zeta$ results in

$$\lim_{l\to\infty} \Pr\left[\min_{1\le k\le l} l_{h,k}(P) > \left(\frac{x(h!)^m}{l(L_{\text{eq}})^{mh}}\right)^{\frac{1}{mh}}\right] = e^{-x}.$$

For large $l = \frac{h!}{z_1! z_2!}$, we obtain the approximate distribution of the minimum length, $l_h = \lim_{l\to\infty} \min_{1\le k\le l} l_{h,k}(P)$, of a h-hop path

$$\Pr[l_h \le y] = 1 - \exp\left(-\frac{h!}{z_1! z_2!}\left(\frac{(L_{\text{eq}} y)^h}{h!}\right)^m\right). \tag{16}$$

The average length of the shortest h path is with $(h!)^{\frac{1}{h}} \approx \frac{h}{e}(2\pi h)^{\frac{1}{2h}} \approx \frac{h}{e}$

$$\begin{aligned} E[l_h] &= \int_0^\infty \Pr[l_h > y]\, dy \\ &= \Gamma\left(1 + \frac{1}{mh}\right) \frac{(h!)^{\frac{1}{h}} \left(\frac{z_1! z_2!}{h!}\right)^{\frac{1}{mh}}}{L_{\text{eq}}} \\ &\approx \frac{h}{e L_{\text{eq}}} \left(\frac{z_1! z_2!}{h!}\right)^{\frac{1}{mh}}. \end{aligned}$$

Since all link weight components are independent and equal in distribution, we can interpret $E[L_{\text{eq}} l_h]$ as the weight of the shortest path in m dimensions. For a square 2d lattice, using ([1, 6.1.49]) $\binom{2z}{z} \approx \frac{2^{2z}}{\sqrt{\pi z}}$, the formula

$$E[L_{\text{eq}} l_h] \approx \frac{h}{e 2^{\frac{1}{m}}} \tag{17}$$

shows that the weight of the shortest path very slowly increases with m as $2^{-\frac{1}{m}}$ and that for any dimension m, $\frac{h}{e2} \le E[L_{\text{eq}} l_h] \le \frac{h}{e}$.

The variance equals

$$\begin{aligned} \text{var}[l_h] &= \int_0^\infty (y - E[l_h])^2 d\Pr[l_h \le y] \\ &= \frac{(h!)^{\frac{2}{h}} \left(\frac{z_1! z_2!}{h!}\right)^{\frac{2}{mh}}}{(L_{\text{eq}})^2} \left(\Gamma\left(1 + \frac{2}{mh}\right) - \Gamma^2\left(1 + \frac{1}{mh}\right)\right). \end{aligned}$$

For large h, we see that

$$\Gamma\left(1 + \frac{2}{mh}\right) - \Gamma^2\left(1 + \frac{1}{mh}\right) = \frac{\pi^2}{6}\frac{1}{(mh)^2} + O\left(\frac{1}{(mh)^3}\right).$$

Hence,

$$\text{var}[l_h] \approx \frac{\pi^2}{6}\frac{(E[l_h(P)])^2}{(mh)^2} \to \frac{\pi^2}{6}\frac{1}{e m^2 L_{\text{eq}}}$$

which is rather small and independent of h as is common for extremal distributions.

C. *Perfect Negative Correlation* $(m = 2)$

In case of $m = 2$ and perfect negative correlation, the first path weight is $W_{h,1} = \sum_{j=1}^{h} u_j$ and the second is $W_{h,2} = h - \sum_{j=0}^{h} u_j = h - W_{h,1}$. Then

$$l_h(P) = \max\left[\frac{W_{h,1}}{L_1}, \frac{W_{h,2}}{L_2}\right] = \max\left[\frac{W_{h,1}}{L_1}, \frac{h - W_{h,1}}{L_2}\right].$$

If $L_1 = L_2 = L_{\text{eq}}$, then $L_{\text{eq}} l_h(P) \ge \frac{h}{2}$ and if $W_{h,1} \le x \le \frac{h}{2}$, then $L_{\text{eq}} l_h(P) \ge h - x$ else $\frac{h}{2} \le L_{\text{eq}} l_h(P) \le x$. Thus, $\Pr[\frac{h}{2} \le L_{\text{eq}} l_h(P) \le z]$ equals

$$\begin{aligned} &\Pr\left[\frac{h}{2} \le W_{h,1} \le z\right] + \Pr\left[h - z \le W_{h,1} \le \frac{h}{2}\right] \\ &= F(z) - F\left(\frac{h}{2}\right) + F\left(\frac{h}{2}\right) - F(h - z) \\ &= F(z) - F(h - z). \end{aligned}$$

Assuming as before independence of paths, then for the minimum length path holds

$$\Pr\left[\frac{h}{2} \le \min_{1\le k\le l} L_{\text{eq}} l_{h,k}(P) \le z\right] = 1 - \prod_l \Pr[L_{\text{eq}} l_h(P) > z].$$

With $\Pr[L_{\text{eq}} l_h] = \Pr[\frac{h}{2} \le \min_{1\le k\le l} L_{\text{eq}} l_{h,k}(P) \le z_l]$,

$$\begin{aligned} 1 - \Pr[L_{\text{eq}} l_h] &= \exp[l \log(1 - [F(z_l) - F(h - z_l)])] \\ &= \exp[-l[F(z_l) - F(h - z_h)]] \\ &\quad \times (1 + o[hF(z_l) - F(h - z_l)]). \end{aligned}$$

If $\lim_{l\to\infty} l[F(z_l) - F(h - z_l)] = \xi$, then $1 - \Pr[\frac{h}{2} \le L_{\text{eq}} l_h \le z_l] = e^{-\xi}$. It remains to find z_l in terms of ξ. We rewrite $z_l = \frac{h}{2} + x_l$. For small x_l and with $f(x) = \frac{dF(x)}{dx}$

$$\begin{aligned} \xi &= l\left[F\left(\frac{h}{2} + x_l\right) - F\left(\frac{h}{2} - x_l\right)\right] \\ &= l\left[F\left(\frac{h}{2}\right) + f\left(\frac{h}{2}\right) x_l \right. \\ &\quad \left. + -F\left(\frac{h}{2}\right) + f\left(\frac{h}{2}\right) x_l + O(x_l^3)\right] \\ &= 2lf\left(\frac{h}{2}\right) x_l + o\left(x_l^3\right) \end{aligned}$$

such that, with the Gaussian approximation for $f(\frac{h}{2}) \simeq \frac{1}{\sqrt{\frac{\pi h}{6}}}$ and $l = \frac{h!}{z_1! z_2!}$

$$x_h = \frac{\xi}{2lf\left(\frac{h}{2}\right)} = \frac{\xi(z_1! z_2!)}{2h!}\sqrt{\frac{\pi h}{6}}.$$

Finally

$$\Pr\left[\frac{h}{2} \le L_{\text{eq}} l_h \le \frac{h}{2} + y\right] = 1 - \exp\left(-2\frac{h!}{z_1! z_2! \sqrt{\frac{\pi h}{6}}} y\right) \tag{18}$$

from which

$$E[L_{\rm eq}l_h] = \frac{h}{2} + \int_0^\infty \exp\left(-2\frac{h!}{z_1!z_2!\sqrt{\frac{\pi h}{6}}}y\right)dy$$
$$= \frac{h}{2} + \frac{(z_1!z_2!)}{2h!}\sqrt{\frac{\pi h}{6}}. \quad (19)$$

Hence, for large h, the average $E[L_{\rm eq}l_h]$ rapidly tends to $\frac{h}{2}$, as has been verified through simulations.

REFERENCES

[1] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions*. New York: Dover, 1968.

[2] M. A. Berger, *An Introduction to Probability and Stochastic Processes*. New York: Springer-Verlag, 1993, Springer Texts in Statistics.

[3] B. Bollobas, *Random Graphs*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 2001.

[4] P. Cheeseman, B. Kanefsky, and W. M. Taylor, "Where the really hard problems are," in *Proc. IJCAI-91*, San Mateo, CA, 1991, pp. 331–337.

[5] S. Chen and K. Nahrstedt, "On finding multi-constrained paths," in *Proc. ICC'98*, vol. 2, 1998, pp. 874–879.

[6] H. De Neve and P. Van Mieghem, "TAMCRA: A tunable accuracy multiple constraints routing algorithm," *Comput. Commun.*, vol. 23, pp. 667–679, 2000.

[7] P. E. Dunne, A. Gibbons, and M. Zito, "Complexity-theoretic models of phase transitions in search problems," *Theoretical Comput. Sci.*, no. 249, pp. 243–263, 2000.

[8] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*. San Francisco, CA: Freeman, 1979.

[9] ——, "'Strong' NP-completeness results: Motivation, examples, and implications," *J. ACM*, vol. 25, no. 3, pp. 499–508, Jul. 1978.

[10] I. P. Gent and T. Walsh, "Analysis of heuristic for number partitioning," *Comput. Intell.*, vol. 14, no. 3, pp. 430–452, 1998.

[11] G. H. Hardy, J. E. Littlewood, and G. Polya, *Inequalities*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 1973.

[12] B. Hayes, "Can't get no satisfaction," *Amer. Scientist*, vol. 85, no. 2, pp. 108–112, Mar.–Apr. 1997.

[13] G. Istrate, "Computational complexity and phase transitions," in *Proc. 15th Annu. Conf. Computational Complexity*, Florence, Italy, Jul. 2000, pp. 104–115.

[14] J. M. Jaffe, "Algorithms for finding paths with multiple constraints," *Networks*, vol. 14, pp. 95–116, 1984.

[15] F. A. Kuipers, T. Korkmaz, M. Krunz, and P. Van Mieghem, "Overview of constraint-based path selection algorithms for QoS routing," *IEEE Commun. Mag.*, vol. 40, no. 12, pp. 50–55, Dec. 2002.

[16] ——, "Performance evaluation of constraint-based path selection algorithms," *IEEE Network*, vol. 18, no. 5, pp. 16–23, Sep./Oct. 2004.

[17] L. A. Levin, "Average case complete problems," *SIAM J. Comput.*, vol. 15, no. 1, pp. 285–286, 1986.

[18] D. H. Lorenz and A. Orda, "QoS routing in networks with uncertain parameters," *IEEE/ACM Trans. Netw.*, vol. 6, no. 6, pp. 768–778, Dec. 1998.

[19] Q. Ma and P. Steenkiste, "Quality-of-service routing with performance guarantees," in *Proc. 4th Int. IFIP Workshop on Quality of Service (IWQoS)*, New York, NY, May 1997, pp. 115–126.

[20] R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, and L. Troyansky, "Determining computational complexity from characteristic 'phase transitions'," *Nature*, vol. 400, pp. 133–137, Jul. 8, 1999.

[21] E. F. Mykytka and C.-Y. Cheng, "Generating correlated random variates based on an analogy between correlation and force," in *Proc. 1994 Winter Simulation Conf.*, 1994, pp. 1413–1416.

[22] D. Pisinger, "Algorithms for Knapsack Problems," Ph.D. dissertation, Dept. Comput. Sci., Univ. of Copenhagen, Denmark, Feb. 1995.

[23] P. Van Mieghem, H. De Neve, and F. A. Kuipers, "Hop-by-hop quality of service routing," *Comput. Netw.*, vol. 37, no. 3–4, pp. 407–423, 2001.

[24] P. Van Mieghem, "Paths in the simple Random graph and the Waxman graph," *Prob. Eng. Inf. Sci. (PEIS)*, vol. 15, pp. 535–555, 2001.

[25] P. Van Mieghem and F. A. Kuipers, "On the complexity of QoS routing," *Comput. Commun., Special Issue on QofIS'01*, vol. 26, no. 4, pp. 376–387, Mar. 2003.

[26] ——, "Concepts of exact quality of service algorithms," *IEEE/ACM Trans. Netw.*, vol. 12, no. 5, pp. 851–864, Oct. 2004.

[27] Z. Wang and J. Crowcroft, "Quality-of-service routing for supporting multimedia applications," in *IEEE J. Sel. Areas Commun.*, vol. 14, Sep. 1996, pp. 1228–1234.

[28] Z. Wang, "On the complexity of quality of service routing," *Inf. Process. Lett.*, vol. 69, pp. 111–114, 1999.

Fernando A. Kuipers received the M.Sc. degree and Ph.D. degree (*cum laude*) from the Delft University of Technology, Delft, The Netherlands, in 2000 and 2004, respectively.

He was a member of an interdisciplinary research center on the Design and Management of Infrastructures, where he participated in the Telecommunications project. Currently, he is an Assistant Professor in the Network Architectures and Services group at Delft University of Technology. His work mainly focuses on the algorithmic aspects, complexity, and dynamics of quality of service (QoS) routing.

Piet F. A. Van Mieghem (M'05) received the M.Sc. and Ph.D. degrees in electrical engineering from the K. U. Leuven, Belgium, in 1987 and 1991, respectively.

He is a Professor at the Delft University of Technology, Delft, The Netherlands, where he is Chairman of the unit Network Architectures and Services (NAS). His main research interests lie in new Internet-like architectures for future, broadband and QoS-aware networks, and in the modeling and performance analysis of network behavior. Before joining Delft, he worked at the Interuniversity Micro Electronic Center (IMEC) from 1987 to 1991. From 1992 to 1993, he was a visiting scientist at the Massachusetts Institute of Technology (MIT), Cambridge. During 1993 to 1998, he was a member of the Alcatel Corporate Research Center in Antwerp, where he was engaged in performance analysis of ATM systems and in network architectural concepts of both ATM networks (PNNI) and the Internet.